



Optical Guidance for Shoreline-Following UAVs

B. Kamgar-Parsi,¹ P. Baker,¹ A. Kahn,² and J. Kellogg²

¹Information Technology Division

²Tactical Electronic Warfare Division

Recent counterinsurgency combat operations in and around inland waters indicate the need for specialized technological capabilities to meet the unique challenges of warfighting in riverine theaters. The ability to conduct close surveillance of enemy activities in the riverine arena is of crucial importance to our warfighters. Over the past several years, the Information Technology Division, in collaboration with the Tactical Electronic Warfare Division, has been developing autonomous unmanned air vehicles (UAVs) capable of guiding their flight course along riverbanks and shorelines. Autonomous UAV guidance based on sensory input rather than GPS promises to increase the flexibility in tasking for single or multiple UAVs. This article describes a guidance and autopilot system that successfully navigated a curved shoreline. The project demonstrated that the near infrared spectrum is particularly good for navigating on shorelines since the land and water are easily distinguishable in these wavelengths, as are objects in the water. Also presented is a new control technique which is able to find control parameters in the image space. This overcomes the need for specialized edge extraction and path generation. Attitude control was maintained by the use of matched long wave infrared thermopiles.

UAVS WITHOUT GPS

Surveillance of terrain, bathymetry, and enemy activities in riverine combat theaters is of crucial importance to our warfighters. To meet the special challenges of operating in riverine and littoral environments, the Information Technology Division, in collaboration with the Tactical Electronic Warfare Division and as a part of the CoastWatcher program, has developed autonomous unmanned air vehicles (UAVs) capable of guiding their flight course along riverbanks and shorelines. Field tests demonstrate that these UAVs can navigate autonomously along waterways and can detect the presence of objects in the water.

Current UAV systems operate by using global positioning system (GPS) waypoints, or are controlled by a direct video downlink (teleoperated), or both. However, GPS may be denied or may be ineffective if maps are not available or are outdated due to the dynamic nature of the riverine environment. Dedicating precious manpower to pilot UAVs is expensive, and the benefits of UAV reconnaissance will not be fully realized until they are autonomous. It is highly desirable to control a UAV without having to specify coordinates and without a remote pilot having to fly it. Furthermore, to conduct surveillance in a riverine theater, it is essential for the UAV to be able to recognize water from non-water. Our UAV guidance system integrates navigation and water detection, simplifying the design and increasing the flexibility and robustness of the UAVs.

We have developed an optical UAV guidance system that directs a UAV along a shoreline via images

captured from a silicon CMOS camera mounted underneath the vehicle. This type of autonomous guidance system can reduce the cognitive load on the warfighter, as well as confer immunity to GPS outages. Autonomous guidance will free up warfighters to give supervisory commands to the UAV, which will report back when a surveillance update is available.

GUIDANCE APPROACH

In order for the UAV to follow a shoreline, it must be able to distinguish water and land, and control itself to remain on or near that boundary. This section describes the general approach to the system, and later sections provide the details.

The system first must identify pixels in the images from the camera as water or land, and put the result into a coordinate system from which the air vehicle can calculate guidance. Because of weight and power constraints, it is important to solve the water detection problem with as little sensing and computation effort as possible. Land and water, while difficult to distinguish in the visible spectrum, are easily distinguished in the near-infrared (NIR) spectrum.¹ Fortunately, the silicon used for conventional image sensors is strongly sensitive to this NIR radiation (700–1000 nm), which is why most consumer cameras come with an infrared-cut filter blocking this “nuisance” radiation. Off-the-shelf cameras with appropriate filters can be used to detect water via a NIR image of the ground.

Since the camera is mounted rigidly on the fuselage of the air vehicle, the image of water versus land is pro-

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007	
4. TITLE AND SUBTITLE Optical Guidance for Shoreline-Following UAVs				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, 4555 Overlook Avenue SW, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

jected onto the ground plane using altitude and attitude information received from the autopilot. Appropriate geometric transforms are applied to accomplish this. There has been prior work on path-following with UAVs, in which a path is first extracted and then a control algorithm directs the aircraft to follow that path.^{2,3} The work described here differs from previous work in that the curve corresponding to the shoreline is not extracted into parametric form. Instead, the algorithm transforms the image into a coordinate system with one axis being a hypothesized bank angle and the other axis being the arc length flown. From here the bank angle resulting in a path tangent to the shoreline can be extracted. This allows the image processing to forgo complicated computations of path, and directly compute the roll angle that will result in the aircraft flying tangent to the path. The simplicity of this algorithm makes it more robust to image processing problems where more complicated edge detectors might fail.

Interestingly, the sensors the UAV uses for attitude control use infrared light, but in a different band, the thermal band. The fact that the sky is colder than the ground allows the attitude to be measured by assuming that the plane is inside a sphere where the top hemisphere is cold and the bottom hemisphere is hot. These sensors are inexpensive and reliable and allow the vehicle to be controlled effectively.

NEAR-INFRARED IMAGING FOR WATER DETECTION

While water has a distinct color and texture, these attributes are highly variable. Detecting water using these cues with a standard red-green-blue (RGB) camera can be difficult, since water can take on many colors, depending on turbidity, depth, waves, and time of day. Simulations using such cues were not sufficiently stable or general. However, NIR imaging worked much better in our simulations, producing images that were much more easily interpreted. The satellite mapping community has used these wavelengths for years for water delineation. Both clear and turbid waters strongly absorb NIR wavelengths from 750 nm to 1000 nm, while most other natural substances (vegetation, soil, rocks) reflect them. Because silicon imagers are quite sensitive to these NIR wavelengths (in fact, more sensitive than to visible wavelengths), the systems can use a standard silicon CCD or CMOS camera with a filter that blocks out wavelengths shorter than 800 nm, although the particular wavelength cutoff is not crucial.

The images, when filtered to capture only NIR wavelengths, are straightforward to process into a map of land versus water. The camera exposure is set so that

the land pixels are bright but not saturated. Given this, the pixel values follow a roughly trimodal distribution. The darkest pixels are water, because they do not reflect much NIR light, and the bright pixels are land. The saturated pixels are reflected sunlight from water. Some of the pixels around the reflected Sun on water will be bright but not saturated, so that some additional image morphology operators need to be applied to remove an annulus around the saturated Sun to improve detection performance. Figure 1 is an example of imagery taken on a sunny day. The water detection algorithm detects water as large areas of uniformly dark pixels, or saturated pixels, and land is everything else. This procedure has proved reliable in live tests, and is a fast operation even on the relatively slow CPU currently installed on the vehicle. The most time-consuming task for the computer is warping the image onto the ground's coordinate system.

AIR VEHICLE ROLL CONTROL

The challenge for the guidance system is to follow an irregular shoreline in real-world conditions. Shorelines are usually highly variable, with multiple inlets and peninsulas, resulting in flight paths that will double back, breaking assumptions of many path-following algorithms. In addition, most autopilots work by maintaining the heading of the airplane, and this is the approach that the system used in the initial simulations. However, this approach was highly ineffective at controlling the airplane in simulation because the vehicle's path is controlled indirectly through heading updates that are in turn controlled by bank angles. It is simpler and more effective to use the guidance to control the airplane's path directly by issuing roll commands.

The guidance system sends the roll angle to the autopilot, which operates in coordinated flight, meaning the rudder operates to keep the airplane at a fixed altitude. For robustness it is desirable for the algorithm to minimize the state information carried from one iteration to the next. This lack of state information means that it is important to keep the coastline in view as much as possible, which means that the airplane should turn around sooner if there is a peninsula that will cause a turn of more than ninety degrees. The algorithm in the work of Frezza⁴ satisfies these constraints, but relies on previously extracted path data, which can be expensive to extract and parameterize. Our algorithm avoids this. The algorithm warps the image into the control coordinate space. This "turn map" has as its x-axis the arc length that the airplane travels, and has as its y-axis possible roll angles. Each pixel thus corresponds to the terrain expected to be underneath the airplane given a fixed roll for a certain distance. Horizontal lines in the turn map correspond to circular



FIGURE 1
In-flight imagery on a sunny day.

paths in the map of the ground plane. Figure 2 shows an actual calculated world map from an experiment, and Fig. 3 shows the turn map generated from that world map. The intuition behind the turn map is that it maps out the predicted terrain under the aircraft for all constant roll parameters, which in the absence of wind, are circles. In most cases, only one of these circles will be tangent to the coastline, and this will occur when there is a sharp transition of terrain between two rows (i.e. possible roll commands).

In this space, the calculation of a suitable path is transformed into a simple image-processing problem of finding a large change from horizontal line to line. The line at which the largest change occurs is the roll angle chosen, thus obviating the need for any fragile edge extraction/parameterization. Many possible trajectories will cause the airplane to follow a trajectory close to the shoreline. The algorithm must find only one of these reasonable trajectories, in the absence of further constraints on the flight path.

FIGURE 2
World example showing separation of water and land.

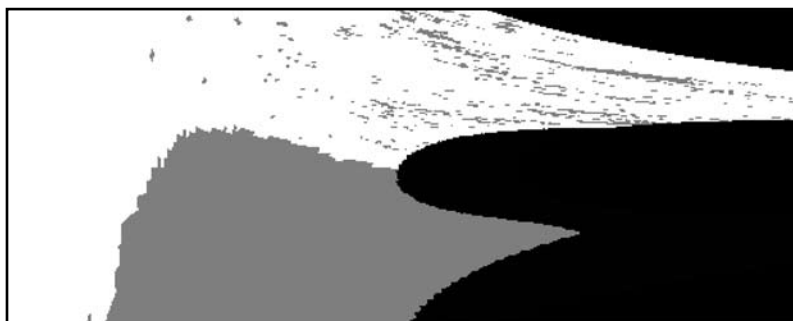
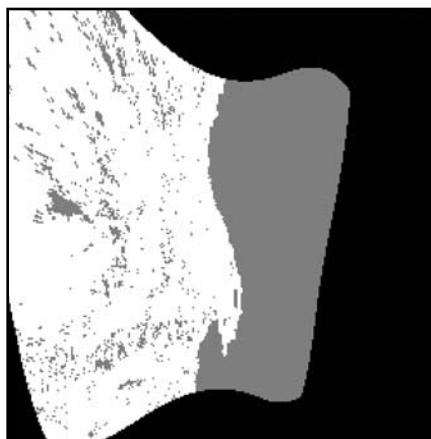


FIGURE 3
Turn map generated from world map of Fig. 2.

The guidance system sends roll commands to the autopilot in order to guide the airplane along the coast. The following algorithm generates the roll commands:

- Analyze the image to find land and water pixels.
- Project image pixels onto the ground plane.
- Compute the turn map.
- Choose the roll angle that corresponds to the largest line-to-line change in the turn map.

Convergence properties of this algorithm are not known at this point, but limited testing has shown it to be a stable method.

GUIDANCE HARDWARE DESIGN

The camera system chosen is a Unibrain Fire-I camera with a resolution of 640×480 pixels. This camera is small and has a Firewire interface, which is easy to program in the Linux operating system. The CCD quality is not of prime concern because the algorithm looks at spectral data, and coarse resolution is adequate for UAV guidance. Using a 1.9 mm focal length lens on the type 1/4 CCD results in a horizontal field of view of 87 degrees and a vertical field of view of 71 degrees. Since the camera is attached to the fuselage of the UAV, a roll of 30 degrees will still result in a minimum 30 degree field of view to each side, sufficient to maintain view of the shoreline in expected flight circumstances.

An Ampro CoreModule 600 PC-104 board runs the operating system and guidance code, while an Ampro MiniModule 1394 connects to the Firewire camera interface. The processor is an Intel ULV Celeron, running at 400 MHz. This system is capable of handling the computational load of the guidance system. Storage consists of a SanDisk 4GB CompactFlash card, ample storage for image logging.

VEHICLE CONTROL

A prime requirement in the design of the autopilot for our work under the CoastWatcher program is to provide true attitude information, preferably without having to rely on the data from a GPS receiver. A simple, yet effective way of providing this attitude information is with the use of infrared thermopiles.⁵ In addition to the thermopiles, a GPS receiver was added to provide a ground-truth reference as to the actual flight path that the aircraft took, for later data analysis. The GPS altitude data was also used in the autopilot for altitude hold functionality, but this same feature could have been implemented as a barometric pressure system. A hobby model aircraft radio control receiver was integrated into the system for testing, launch, and recovery of the vehicle. The autopilot system commu-

nicates with the guidance computer via RS-232 serial. A system-level diagram of the autopilot is shown in Fig. 4.

The vehicle keeps no state and computes new roll angles at 5 Hz. The algorithm has proved to be stable and smooth in simulation and operation.

IR SENSOR OPERATION

The attitude measurement system of the autopilot is based on the use of MEMS thermopiles. These devices convert incident long-wave infrared radiation into a proportional voltage. When measuring attitude, these sensors sense the average background temperature of the sky and ground visible within each sensor's field of view. In most cases, there is a significant difference in temperature between sky and ground. During testing, a total temperature differential of up to 60 °F has been measured. An illustration of the roll sensing is shown in Fig. 5. As can be seen, the sensors are arranged such that a differential temperature between the port and starboard side of the aircraft can be measured.

This differential temperature, or voltage measurement, can then be used to compute the vehicle attitude relative to the thermal horizon. If the vehicle is flying at a high altitude, such that the ground can be assumed flat over the field of view of the sensor, then the assumption is valid in that the thermal horizon is equal to the inertial horizon.

In implementation, three pairs of sensors are used. The sensors are built into the structure of the vehicle, arranged to minimize the amount of aircraft structure in each sensor's field of view; this maximizes the accuracy of the attitude measurement. The sensors are arranged such that one pair is port and starboard to measure the roll attitude. Another pair of sensors is arranged fore and aft to measure pitch attitude. A final pair of sensors is arranged top and bottom to measure the maximum thermal differential possible. This last vertical pair is needed so that an accurate calculation of attitude can be made.

CONTROLLER

The controller design for the autopilot is based around three single-input single-output (SISO) linear controllers to regulate roll attitude, pitch attitude, and altitude. It is assumed that the base air vehicle is statically stable, and is highly damped in both roll and pitch, thus not requiring additional rate damping. The control design is also made around a single operating point. As the autopilot design does not include an airspeed sensor, no gain scheduling is possible to handle the change in sensitivity of the control surfaces with airspeed.

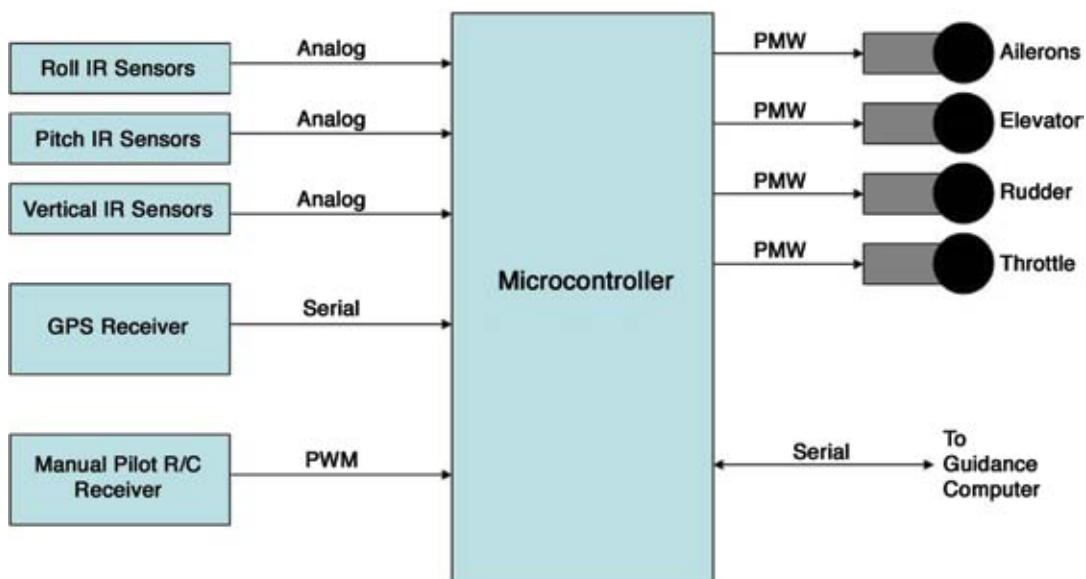


FIGURE 4
Autopilot architecture.

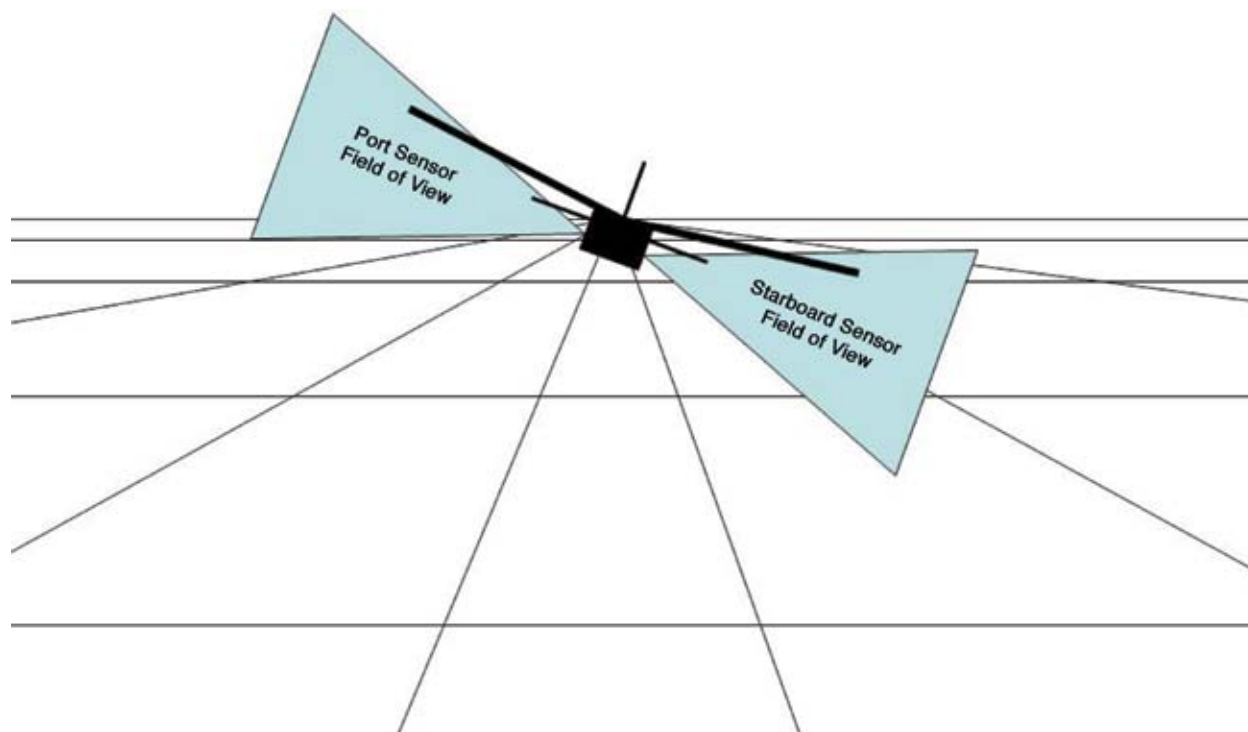


FIGURE 5
Roll sensor fields of view.

The roll and pitch control loops are both implemented as proportional-integral controllers featuring integrator anti-windup protection. The altitude control loop is a proportional-integral controller with flight path angle command as a feed-forward into the pitch control loop. Altitude measurement and commanded altitude are used with the gain based on the desired convergence time and designed flight speed to compute a bounded flight path angle. The resulting flight path angle is then used as the command for the pitch attitude controller.

VEHICLE DESCRIPTION

The airplane is roughly five and a half feet long and has a seven-foot wingspan. For propulsion it has an electric motor connected to two lithium polymer cells, which provide greater than thirty minutes of flight time when fully charged. The airplane has a GPS antenna, the output of which is used for altitude control, and to provide a ground track with which to judge the success of the coast following. The vehicle also has a FreeWave modem for downloading telemetry during the flight, although all data is logged as well.

TESTING AT BLOSSOM POINT PROVING GROUNDS

Our team took the UAV to Blossom Point Proving Grounds, MD, and flew the airplane on numerous tests to gauge its water-following abilities at different locations, including at the sharp peninsula depicted in Fig. 6. The aircraft took off under radio control, shown in Fig. 7. When the vehicle achieved an altitude of 200 m and was headed toward the coast, control was handed over to the guidance and autopilot system. The airplane consistently acquired the shoreline and rounded the peninsular point in both directions, maintaining altitude. This demonstrates the capacity to detect water and land, convert this information to roll commands, and guide the vehicle so that it stays close to the shoreline.

In Fig. 6, the track of the airplane as obtained by the GPS antenna is superimposed on a map of Blossom Point. The guidance system was engaged at the green end of the track, flying directly toward the shoreline. The track colors indicate

- red: hard right turn
- yellow: gentle right turn
- green: approximately straight
- cyan: gentle left turn
- blue: hard left turn

From the track you can see that the guidance system started while the aircraft was over land. The aircraft then acquired the shoreline, rounded the point,



FIGURE 6

Aircraft track at Blossom Point Proving Grounds. The track has not been corrected for wind effects, which can be accomplished with a simple wind module.



FIGURE 7

Aircraft used in testing.

and continued somewhat off the shoreline until control was retaken by the pilot.

The testing day was sunny, with a south wind of approximately 12 m/s, a significant fraction of the vehicle's airspeed. The wind estimates were obtained by using the aircraft's constant airspeed and the GPS coordinates, and finding the center of a circle fitted to the ground velocity estimates. The airplane tracked the shoreline in both instances without oscillations.

The northward bias in the track is due to the wind. The airspeed of the plane was approximately 30 m/s, and when the vehicle is traveling east to west, the bias is between 100 and 200 meters. The vehicle looks ahead a maximum of 500 meters. The wind speed of 12 m/s caused the aircraft to crab at about 20 degrees. Since the algorithm, when following a straight line, causes the airplane to aim for the farthest land point, this crab angle corresponds to approximately the bias seen. The bias was seen in both directions on both the north and

south coasts of the point, so it was not due to image processing.

CONCLUSION

This project has shown that an optical guidance system can control a small UAV. The test lasted for approximately three minutes, but the guidance system could have kept guiding the airplane for a longer time, had the radio control had a long enough range. The UAV acquired the shoreline, turned appropriately, and continued without oscillations around a peninsula.

The use of infrared thermopiles provides a reliable, simple, and inexpensive method for the autopilot to measure the attitude of the airplane. Alternative attitude control systems using gyroscopes are heavier and can have problems with integration error over long time periods. The success of this aircraft's flight shows that this technology can form an important attitude sensor in the future.

The use of infrared spectral information (both near and far) allows the guidance algorithm to be robust and simple, and constitutes an example for the design of autonomous guidance systems for UAVs. There has been much research into the spectral properties of water and land, and this knowledge can be used to navigate autonomously, not only to generate maps.

A UAV in a riverine environment can give crucial advance notice of possible courses, and locations of

threats, by flying ahead of a patrol boat or reporting to warfighters stationed in or near water. Future versions with sensors operating in different wavelengths could conduct bathymetry. This capability, combined with autonomous navigation, could enable a UAV to accomplish tasks such as finding the closest break in the reef so that the boat can land on the shore, or finding the best path for a boat down a river. This research supports the Navy's need for tactical intelligence information in riverine and littoral environments, bringing the surveillance capabilities to the warfighters who need them the most.

[Sponsored by ONR]

References

- ¹ W.R. Phillipson and W.R. Hafker, "Manual versus Digital Landsat Analysis for Delineating River Flooding," *Photogrammetric Engineering and Remote Sensing* **47**, 1351-1356 (1981).
- ² E. Frew, T. McGee, Z. Kim, X. Xiao, S. Jackson, M. Morimoto, S. Rathinam, J. Padial, and R. Sengupta, "Vision-Based Road-Following Using a Small Autonomous Aircraft," *Proc. 2004 IEEE Aerospace Conference* **5**, 3006-3015 (2004).
- ³ S. Rathinam, Z. Kim, and R. Sengupta, "Vision-Based Following of Structures Using an Unmanned Aerial Vehicle," Tech. Rep. UCB-ITS-RR-2006-1, Institute of Transportation Studies, University of California at Berkeley, March 2006.
- ⁴ R. Frezza, "Path Following for Air Vehicles in Coordinated Flight," *Proc. 1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 884-889 (1999).
- ⁵ B. Taylor, C. Bil, S. Watkins, and G. Egan, "Horizon Sensing Attitude Stabilization: A VMC Autopilot," *18th International UAV Systems Conference, Bristol, 2003*, Paper 22 (2003). ★